

# Biologically-Inspired Cognitive Architectures Program

## Questions and Answers – Industry Day January 27, 2005

1. What is the relationship, if any between ACIP and BICA?

*ACIP is oriented toward high performance hardware platforms.*

2. Can different parts of the same company, firewalled, bid on Task C as well as Task A or B?

*Yes. This is not a legal conflict of interest problem. We don't want the same people designing the test problem that their system will be tested against. We do, however, want the people working on Thrust C to be very collaborative with people working on Thrusts A and B. The Thrust C performers will have to interact with Thrust A&B performers to get recommendations on what the test problems should be. Thrust A and Thrust B performers would likely propose different types of tests if it were up to them, so it will be a Task C job to figure out a compromise for what a reasonable test suite would be. So as long as there is reasonable separation, this is fine.*

3. How general does the cognitive architecture have to be?

*It has to be general. One of the goals of this program is the integration we've talked about. We want an architecture that is capable of doing it all. Now we're not necessarily going to build it all, but we do want an overarching architecture that is quite general.*

4. If focus is on a few applications, then may not serve all cognitive examples (decathlon) in Thrust C?

*During this first year, we're really trying to define what the test problems would be for the second phase of the project. For the cognitive decathlon we do want that test to be a general suite of tests that includes perception, memory, reasoning, and other major cognitive functions. We hope to construct the test suite by taking from the collection of experimental psychology tasks where you have evidence of how humans do these cognitive things - that's the spirit. It would be very general in that we would want the same system to be able perform the entire suite of tests, and not allow developers to develop and optimize a solution for each one test individually. And for the applications, they will necessarily be more specific, but again we'd like to come up with a range of those.*

*And you'll notice that we're doing the decathlon in the first half of the program. When we create the challenge problems in the second half we'll*

*want to make sure that we have a suite of challenge problems that exercise all the elements in the decathlon.*

5. Is it sufficient to create a cognitive agent, but one that can't communicate with other cognitive agents?

*No. There would have to be communication in some form in the system. Some of these questions we won't have the answer for yet. For example, what do you do about perception? Raw sensory information? Are we just going to have the context of the application be, "here's raw sensory data, start from scratch and process it?" Or do we abstract that into some interface where you get some symbolic expression of the world that you're dealing with, and then the system processes the symbolic information? The same with language, will it just open ended and your cognitive system will have to handle any natural language or will we have some API where we can find some natural but somewhat limited language set? We don't know. That's what the test developer will have to figure out.*

6. You mentioned the tests a lot. I'm assuming that you are looking for cognitive theories and perhaps additionally cognitive architectures in the proposals. Are you looking also for proposals for how to test or are you talking about tests you will develop to be applied to products?

*Thrust C performers are the people developing the test. However, we are expecting that, as part of your proposals in Thrust A or B, you would be developing ideas on how your system would be tested that would be presented to the test developer. You don't have to present those ideas in your proposal though; we would instead expect the Thrust C developers to be proposing ideas on good tests for cognitive systems. The developers in Thrust C will be the arbitrators, and with us (IPTO) will make the final decision as to what that test really should be.*

7. What if "80% of structures" are not mapped to brain functions?

*Nothing happens. The 80% figure is a goal. We wanted to give you a sense that we are not insisting on 100%, but 50% isn't good enough either*

8. If a team is interested in taking a particular function such as memory and approaching it from both psychological and neurological perspective, which thrusts and tasks should the team apply for?

*We imagine that there would be some people who would want to work from each perspective and have combined approaches. A combined approach on one particular function is OK. Please apply under Thrust A or B or both.*

9. And which tasks? A1 or B1? (from audience)

*You certainly can propose for just A1 or B1, but you also have to propose to Task 2 as well, regardless. Bidders may propose on any combination*

*of A1, B1, A3, or B3 with the following constraints: everyone must bid on Task 2 (A2 or B2;) and any one that bids on Task 3 (A3 or B3) must bid on Task 1 (A1 or B1) as well.*

10. How strongly correlated are you expecting new architectures to be to recent work in cognitive psychology as opposed to classical psychology and traditional symbolic AI?

*To some extent, traditional symbolic AI is what we're trying to get beyond, but there are good ideas in there that you don't want to just abandon. Overall, we certainly want to take advantage of the most recent ideas in cognitive psychology, especially in Thrust A where we're looking at building models that have functional components, such as memory, that mimic how humans behave while performing tasks. You can't have symbolic computations or more traditional software approaches used that don't mimic the way neurobiology works. And we want to account for the most recent results in cognitive psychology. We do want approaches that are based on mimicking the psychology of how human thought works as opposed to purely symbolic AI, which may be logic or some kind of reasoning abstraction that doesn't at all correlate to humans.*

11. What do you expect the role of large-scale, general knowledge to be relative to bottom-up learning?

*We don't want these systems to be ones that require huge amounts of knowledge engineering to put knowledge in the system. We would rather have systems that learn as they go. Given that, even a one year old child has, in computer science terms, an incredible library of common-sense knowledge. We certainly expect some engineering to have to be done up-front, but once that is done you should have a system that would have a fighting chance to learn on its own. It has to be able to add new knowledge, add new procedures, etc. without having to be re-engineered every time there is a new task. We're not looking for something like a new task for a robot sweeping the floor, where you first program 500 sweeping procedures and put that into the system.*

12. In terms of functions: visual, auditory, text-based processings, etc. – what are the emphases or distribution of funds?

*Integration, integration! We want a balanced program. And we're interested in things like affect, emotion -- those kinds of things as well.*

13. What is the priority on technology transition?

*In this case, it is almost none. In this first phase, the goal will be the creation of these architectures which are then ready for Phase 2, where people will go off and implement these architectures and build them to be tested against some problems which then may begin to have more of a transition element. But Phase 1 of this project is much more basic science*

*in its approach, like how these architectures are going to be put together from elements in psychology, computer science, and neuroscience..*

14. Can ACIP contractors submit proposals for BICA?

*Yes.*

15. What is the timeframe for the first set of awards?

*1 July 2005, depending on iterations needed with the DARPA Director. For the purposes of writing proposals, please use 1 July and we'll adjust as needed.*

16. Have you considered a Problem Domain Description Language, such as McDermott's PDDL 2.1 for challenge problems?

*Not specifically, but that would be the kind of thing we'd expect our test problem developers in Thrust C to do. They'll have to develop a test environment. It will have to be an API of some sort that you'll have to plug your cognitive systems into when you start developing them in Phase 2.*

17. If the goal is hybrid architectures – merging traditional AI with neuroscientifically inspired systems that “grow” – how are interfacing issues to be addressed? There may be interesting problems related to the “moving target” nature of learning systems. More specifically, are such interfacing issues to be addressed in Task 1 or Task 3? Is this strictly a Thrust B issue?

*The goal is hybrid architectures. We're not sure how to answer that. We're not really sure it would be a strictly Task 1 or Task 3 or Thrust B issue – but we are definitely interested in systems that grow.*

18. Were there any initial seedling projects related to this BAA, and if so what was the nature of the projects?

*There may not have been any that were specifically part of this one, but there have been some seedling projects that we now see in retrospect were related to the BICA program. There are lots of projects in the area of learning. One of those was with the SOAR people at the University of Michigan where they were trying to add episodic learning to the SOAR model. There was another with John Anderson at CMU who was trying to build up a model in ACT-R where it could simulate an algebra student working through problems in Algebra I. In addition there was also a project with the Neurosciences Institute in La Jolla, CA that was taking their model of the individual processing systems in monkeys and applying it to the application of a soccer-playing robot. The robot is hybrid but the robot uses their visual processing model (Segway platform).*

19. Given the scientific nature of Phase I, do you anticipate teams to be industry/academic, academic-only, or industry-only?

*Any of the above, although we'd certainly expect academic participation so we'd anticipate them to be academic-only or industry/academic. Although, if there are industry research labs that would be interested in this area, they'd be very welcome as well. Since it's more of a research task than an application task, we'd expect the mix to be a little heavier on the academic side.*

20. Are you interested in architectures that include multiple BICA nodes?

*No. Eventually your individual cognitive agent will have to work well with others but we're not interested in building an architecture for cooperative agents. We'd like to have other programs in that area, but this is not one.*

21. Are reinforcement-learning techniques coherent with the program objectives (Thrust A)?

*Absolutely.*

22. How does Thrust C compare with the concept of "Open Experimental Platform (OEP)" that other DARPA programs have used before?

*It doesn't in this phase. In Phase 2 we'll be looking for integrated teams that have all of that.*

23. Is there a list of previously published proposed challenges that would be part of Thrust C?

*Not really. We have at different times done things to generate ideas for grand challenges in AI. I think we have informally accumulated lists of those that we could make available to people who began working on Thrust C. But we don't plan on publishing them on a web page for all to see and we don't want to prejudice the answers by what we've had in the past.*

24. What sort of a final product do you expect to be delivered for Thrust A?

*In the case of Thrust B (neuroscience thrust), it is documentation on those three products that were specified in the PIP. In the case of Thrust A we are expecting the products to be designs of cognitive models that could be implemented in Phase 2. (also see the answer to Question 31.)*

25. What fraction of all listed topics do you expect to be integrated in one proposal? (inaudible explanation from audience)

*I expect a lot of the proposals to be looking at specific aspects of the problem. Some people may choose to focus on the visual system, some people may choose to focus on memory. And when they do that, the completeness criterion applies to that part of the problem. So if you're doing memory, I'm looking for 80% of memory kinds of functions being mapped. The integration across those things I'd like to happen at part of the program during exchanges and evaluations. In general, there is*

*nobody out there that can do it all. We want to bring together people that collectively can cover the problem.*

26. [Follow-up to previous question] If someone could do it all, then is it safe to say that the end thing you're looking for is a hybrid architecture that can grow and that merges AI in the general sense with neurological systems? Is that correct?

*In Thrust B, in the first part of the thrust (Task 1), we want to lay the groundwork for creating the second part – hybrid architectures that are inspired by neuroscience. In the first part, we are creating three kinds of principles that will be used by hybrid design teams in the second part of the program to create hybrid designs using the best of AI along with the principles established in the first part of the program (Task 1). So in the second part of the program, Task 3, those teams ought to have people from neuroscience as well as from conventional AI. The architectures that will be created in Task 3 will need to be complete. But the theoretical work in Task 1 may cover only part of the problem. Those doing Task 3 will have the benefit of all the work done in Task 1.*

27. Will you allow a joint proposal for A and B?

*Yes.*

28. Will you allow a separate proposal for C?

*Yes. But if it's the same company that proposed on or A or B, then different people need to be involved. See the answer to question 2 above.*

29. Please elaborate on the technical exchange task. How should these be budgeted into the proposals?

*There would be one or more meetings, or mini-workshops, where you'd come in and present your ideas from the first 6 months to all the other program participants. The time involved might be something like 2 interchange meetings of 3-4 days each. For purposes of the proposal assume 2 exchange workshops, one in Los Angeles and one in Boston.*

30. Will Phase 2 be open for everyone or Phase 1 winners only?

*It will be open for everyone. We would expect different groups of people for Phase 2 because of the nature of the work.*

31. Reiterate what you envision as a final product from A, B, and C.

*For A and B, it would basically be reports and briefings of your ideas in Task 1. For Task 3, there would be a report that contains your architecture. The idea is that at the end of Phase 1, there would be a report that would be made available to bidders on Phase 2. Below is what we expect to see in the final products for Task 1 and Task 3 for both Thrust A and B:*

*Task 1 Products (Thrusts A&B):*

*General description of theory required? Yes.*

*Descriptive account of principles, computations, mechanisms, and functional, mapping required? Yes.*

*Integrated architecture (80% comprehensive) required? No; but the more comprehensive the better.*

*Explicitly defined algorithms and computations required? No.*

*Code specified? No. Software will be developed in Phase 2.*

*Hardware specified? No. This will be a Phase 2 problem.*

*Task 3 Products (Thrusts A&B):*

*General description of theory required? No (see Task 1).*

*Descriptive account of principles, computations, mechanisms, and functional mapping required? No (see Task 1).*

*Integrated architecture (80% comprehensive) required? Yes.*

*Explicitly defined algorithms and computations required? Yes.*

*Code specified? No. Software will be developed in Phase 2.*

*Hardware specified? No. This will be a Phase 2 problem.*

32. Do you envision a user sitting down and pushing buttons?

*In this program, we are not going to DoD users at all. The users of the product of Task 1 will be the architecture design teams in Task 3. The users of the product of Task 3 will be the applications builders in Phase 2. So we don't have any users in the traditional DoD sense.*

33. Will Thrust C be conducted in a supercomputing environment?

*Probably not. Possibly the tests in Phase 2 might be implemented in a supercomputing environment.*

34. The BAA refers, predominantly, to software solutions. Will a hardware-centric proposal be relevant as well?

*I can't imagine that would be the case, but who knows? One of the goals of the program is to inspire new hardware device creators. To design architectures that would provide inspiration for a new class of devices and new kinds of hardware would be a plus. But this is a software-centric program.*

35. As regards to team composition, do you have a preference for large teams or small teams?

*Medium sized teams. We're not looking for one or two mega teams and we're not looking for dozens of projects exploring niches. In Tasks A1 and B1, there's more room for individual performers. For Tasks A3 and B3, we're looking for more medium sized teams, but individual performers are*

*welcome as well. We want proposals from groups with the right skills and do not have a preconceived notion of team size should be, large or small.*

36. Are notional “FTEs” the total for all efforts or limits for each one?

*Total for all efforts, but very notional. They are rough guesses at best. These numbers are very rough estimates intended to indicate the relative effort between tasks 1 and 3. The total funding for the program will depend primarily on the quality and variety of proposals we receive.*

37. Are you thinking grad student FTEs or researcher FTEs?

*Researcher.*

38. Do you prefer proposals from integrated teams covering many brain components in Task 1 or will integration of various components be done in later tasks?

*Yes. See answers to questions 3, 25, and 26 above.*

39. Do you envision an “integration contractor” for A3 or B3?

*There’s not one single integration contractor for the whole program, rather medium sized teams that can handle the integration of the tasks. However, for Task 1 we would entertain a proposal for someone who wanted to be a “research coordinator” or something like that.*

40. Are you open to meet with us to discuss our ideas? Or to take phone calls?

*Yes, but we’re limited by availability. Please start by emailing questions to the BAA email box [BAA05-18@darpa.mil](mailto:BAA05-18@darpa.mil). From there we could work into phone calls or meetings.*

41. You mentioned CMU, a major AI center, as doing some related work. How do the approaches of AI differ from those that BICA will develop?

*We are interested in biologically inspired ideas –inspired by modeling psychological or neurobiological processes of the human mind. When we mentioned CMU, we were referring in particular to CMU’s history of developing computational models of psychological processes. But, in general, the traditional logic and symbolic reasoning techniques of AI have not achieved the flexibility and ability to deal with the novelty that people can deal with. We want to inject into that framework some new ideas*

42. For Task A, is the goal to exceed human behavior? Is it acceptable to produce systems that “only” match, explain, and predict human behavior? Basically, do you want super-human behavior? Or are psychologically-constrained models/components acceptable?



*We're really after capturing the magic that happens in human cognition that we can't duplicate in computers. If you can match the mistakes in how human memory works, that would be good.*

43. What is the deliverable for Task A1 and B1? Is it executable and validated models? Are descriptive (paper) models/frameworks acceptable?

*Descriptive models and reports are exactly what we're looking for, not code.*

44. What is your interest in distributed or net-centric solutions?

*Not our primary interest.*

45. Ideally, when would you like to see proposals in order to prepare for your briefing to the DARPA Director?

*The primary closing date is 1 March. The BAA will still be open for almost a year beyond that, but the money is based on a 1 March deadline.*

46. How many awards do you expect for Task 1? Task 3?

*It is very open-ended. We are probably expecting to fund 2-6 proposals in each track. Perhaps more in A1 or B1. But we really don't know. It depends on what proposals are submitted.*

47. When do you expect the Phase 2 BAA to be released?

*It's really too far in the future to tell. We'll be basing the pitch for Phase 2 on the results of Phase 1.*

48. What range of sensory inputs do you consider appropriate? Most examples are human-based, but many other sensory mechanisms exist, i.e. RF, IR, infrasonic? Are there constraints?

*Our target is to handle human like cognition. If you can handle vision and sound, that would be enough.*

49. Is the creation of small proof-of-concept capabilities acceptable with Phase 1? If so, is this a natural effort for Task 3?

*We're not developing capabilities in Phase 1, we're developing theories to inspire architectures to be developed in Task 3.*

50. What level of testing do you expect to be done in Task 3?

*None. We're not doing testing in this part of the program at all. We're creating test frameworks. Actual testing will be done in Phase 2.*

51. What range of problems will a single design be expected to address?

*The designs need to be able to do it all. The architectures in Task 3 are general architectures that cross the whole range of capabilities you can have.*

52. True or False: Proposers can propose to specific capabilities, i.e. memory, learning mechanisms, etc. (Tasks 1-3).

*True. But Task 3 does have to be general.*

53. True or False: IPTO wants the aggregate of these thrusts to map to 80% of cognitive functions?

*True.*

54. Can bidders for Thrust C only bid for Tasks 2 and 3?

*Yes.*

55. Can you provide some of your expectations for the team composition and organization composition (i.e. multi-organizational collaboration preferred or is it really just about the ideas and credibility to execute)?

*This is about the ideas. To the extent you need multiple people to have good ideas, that's how big the teams should be.*

56. Can you translate FTE to dollars?

*No. That's up to the proposers.*

57. What do you think of Redwood Neuroscience Institute (RNI)?

*They are very interesting and we'd love to see a proposal from them.*

58. How about an A1 or B1 proposal with an option for A3 (or B3)?

*That would be fine.*

59. Can you imagine the team expanding for Task 3?

*We would expect that to be the case.*

60. Would you entertain pre-proposals, i.e. brief "white papers" summarizing our ideas?

*No.*

61. What again is the end product from Thrust C?

*There will be 2 things, both of which will be frameworks for doing testing in the follow-on part of the program. The first is what we're calling the cognitive decathlon – a set of capabilities that cognitive systems must have. And the other is a set of challenge problems for cognitive systems. At the end of Thrust C we would like to have a specification of the*

*Cognitive Decathlon test suite and the Challenge Problems that could be published in the BAA for Phase 2 of the project.*

62. Given the meetings you see in Task 2, how much of the tests in Thrust C could emerge from those meetings?

*Very much, we'd expect that the test suite would be the result of the negotiation between everyone involved..*